arXiv:math/0611585v3 [math.CO] 27 Jan 2008

Conductance and canonical paths for directed non-lazy walks

Ravi Montenegro [*]

Abstract

We show two Conductance-like theorems for mixing time of non-reversible non-lazy walks; the first holds for walks with small holding probability, while the second theorem holds even for walks with no holding probability. These are used to derive two canonical path theorems for such non-reversible non-lazy walks. As an application we show that a known bound for mixing time of walks on undirected Cayley graphs applies to all finite directed Cayley graphs.

Keywords : Mixing time, canonical paths, evolving sets, Cayley graph.

1 Introduction

The notion of conductance [5] has been key to studying the convergence rate of many random walks. There have been many improvements on this original concept: the method of average conductance showed better mixing bounds for lazy reversible walks if small sets have high conductance [7], evolving sets were used to apply this to lazy non-reversible walks [9], modified conductance extends this to non-lazy non-reversible walks [8], and blocking conductance shows improved bounds for a lazy reversible walk when ergodic flow is not heavily concentrated on a few vertices [6].

In this paper we show a mixing time bound which encompasses all of these concepts. This is done by using evolving sets to study general finite Markov chains in a fashion resembling blocking conductance. In particular, when measuring ergodic flow from a set we count flow into each vertex only up to some ratio r of its size; if the flow is well distributed among vertices then this will differ little from regular ergodic flow even for fairly small values of the ratio r. Two versions of our bound are shown, when there is a non-zero holding probability at every vertex or when there is not, although up to a constant factor the latter bound is stronger than the former.

These new results are then used to show two canonical path theorems for general finite Markov chains. The first requires only a small non-zero holding probability at each vertex, while the latter requires none. This does away with the largely artificial assumption of strong aperiodicity or reversibility required for most canonical path analysis, conditions which are undesirable as algorithms are almost never implemented with a large holding probability and as a non-reversible walk may converge faster than a reversible walk. Although not directly comparable to past bounds, our result says roughly that the Poincaré approach to bounding mixing times [2, 3, 10] applies in the general setting if the maximum path length is replaced by a notion of maximum vertex congestion. Average path length and average vertex-congestion are equal (see Remark 4.6), and although it is usually easier to keep path length small than to smooth the paths among all the vertices, when looking at a

---

[*] Department of Mathematical Sciences, University of Massachusetts Lowell, Lowell, MA 01854, ravi_montenegro@uml.edu.

non-reversible walk it may be worth the extra effort to attempt this so that the artificial assumption of a large holding probability can be dropped.

As an application of sorts, we show that known complexity results for the (non-lazy) simple random walk on an undirected graph with self-loops apply just as well to the simple random walk on an Eulerian graph with self-loops (i.e. strongly connected with in-degree=out-degree at each vertex), and often even without self-loops. Other canonical path methods result in the loss of a factor of max-degree d. A more interesting problem is to study random walks on Cayley graphs (i.e. on groups), for which we show that known bounds [1, 2] for a lazy walk or a walk with a symmetric set of generators can be extended to the non-lazy non-symmetric case.

Our paper proceeds as follows. We review past methods in Section 2. In Section 3 we show our conductance-like mixing results when ergodic flow is spread over many vertices. Then in Section 4 canonical paths are used to bound mixing time in terms of notions of edge and vertex-congestion. We finish with a few examples in Section 5.

2 Preliminaries

Throughout the paper we consider a finite ergodic Markov kernel P (i.e. transition probability matrix) on state space V with stationary distribution π. The walk has minimal holding probability α if $\forall v \in V : \mathsf{P}(v,v) \geq \alpha$, is called lazy if we may take $\alpha \geq 1/2$, and is reversible if $\mathsf{P}^* = \mathsf{P}$ where the time-reversal $\mathsf{P}^*(x,y) = \frac{\pi(y)\mathsf{P}(y,x)}{\pi(x)}$. The ergodic flow from $A \subset V$ to $B \subset V$ is $\mathsf{Q}(A,B) = \sum_{x\in A, y\in B} \pi(x)\mathsf{P}(x,y)$. The chi-square distance between distributions σ and π is

$$\left\|\frac{\sigma}{\pi} - 1\right\|_{2,\pi}^2 = \sum_{v\in V} \pi(v) \left(\frac{\sigma(v)}{\pi(v)} - 1\right)^2 .$$

The L^2 mixing time $\tau_x(\epsilon) = \min\{n \mid \|k_n^x - 1\|_{2,\pi} \leq \epsilon\}$ for density $k_n^x(y) = \frac{\mathsf{P}^n(x,y)}{\pi(y)}$ is the time for chi-square distance to drop to ϵ^2 when started at state $x \in V$.

In order to understand what is new about our results it is necessary to review earlier work on conductance and canonical path methods.

2.1 Conductance methods

Jerrum and Sinclair [5] used an argument involving spectral gap to show that mixing time of a lazy reversible walk can be bounded in terms of the conductance $\Phi = \min_{A\subset V} \Phi(A)$ where $\Phi(A) = \frac{\mathsf{Q}(A,A^c)}{\min\{\pi(A),\pi(A^c)\}}$. The method of spectral profile [4] can be used to extend this to consider set sizes for even non-reversible non-lazy walks, with

$$\tau_x(\epsilon) \leq \int_{4\pi(x)}^{4/\epsilon^2} \frac{2\,ds}{\frac{\alpha}{1-\alpha} s\Phi^2(s)}$$

where the *conductance profile* is given by $\Phi(s) = \min_{\pi(A)\in(0,s]} \Phi(A)$ if set size $s \in (0, 1/2]$, and $\Phi(s) = \Phi(1/2)$ if $s > 1/2$. When there is no holding probability then a few approaches are possible (see [4] and Section 4.4 of [8]), with

$$\tau_x(\epsilon) \leq \int_{4\pi(x)}^{4/\epsilon^2} \frac{4\,ds}{s\Phi_{\mathsf{PP}^*}^2(s)} \qquad \text{and} \qquad \tau_x(\epsilon) \leq \int_{4\pi(x)}^{4/\epsilon^2} \frac{2\,ds}{s\tilde{\phi}^2(s)}$$

where $\Phi_{\mathsf{PP}^*}(s)$ is conductance profile of PP^*, and the *modified conductance profile* is

$$\tilde{\phi}(s) = \inf_{\pi(A)\in[0,s]} \tilde{\phi}(A) \qquad \text{where} \qquad \tilde{\phi}(A) = \frac{\Psi(A)}{\pi(A)\pi(A^c)}$$

and

$$\Psi(A) = \min_{\substack{B\subset V,\, v\in V,\\ \pi(B)\le\pi(A^c)<\pi(B\cup v)}} \mathsf{Q}(A,B) + \frac{\pi(A^c)-\pi(B)}{\pi(v)}\mathsf{Q}(A,v)\,.$$

This is the ergodic flow from A to the worst set of size $\pi(A^c)$, in contrast to $\mathsf{Q}(A,A^c)$ which is the ergodic flow to a specific set of size $\pi(A^c)$. In Lemma 4.19 of [8] it is shown that $\tilde{\phi}(A) \ge \frac{1}{2}\Phi_{\mathsf{PP}^*}(A)$, so up to a small constant modified conductance is the better method.

2.2 Canonical Paths

Before discussing canonical paths results we require some notation:

Definition 2.1. For every $x,\, y \in V$, $x \neq y$, define a path γ_{xy} from x to y along edges of P, and let Γ be the collection of these paths. The *edge-congestion* is given by

$$\rho_e = \rho_e^{\mathsf{P}}(\Gamma) = \max_{a,b\in V:\mathsf{P}(a,b)>0} \frac{1}{\pi(a)\mathsf{P}(a,b)} \sum_{\gamma_{xy}\ni(a,b)} \pi(x)\pi(y)$$

The *maximum path length* is $\ell = \max_{\gamma_{xy}\in\Gamma} |\gamma_{xy}|$.

Canonical paths are used to bound mixing time by considering edge-congestion alone [5], or through a Poincaré inequality involving maximum path length as well [3, 10, 2]. Combining these results in a bound for general Markov chains with non-trivial holding probability α:

$$\tau_x(\epsilon) \le \frac{\rho_e}{2\alpha}\, \min\{4\rho_e,\, \ell\} \log\frac{1}{\epsilon\sqrt{\pi(x)}} \tag{2.1}$$

The Poincaré case can be extended to non-lazy reversible walks [2, 3] if Γ' contains only paths of odd length and also has a path from each vertex x to itself:

$$\tau_x(\epsilon) \le 2\,\rho_e(\Gamma')\ell \log\frac{1}{\epsilon\sqrt{\pi(x)}} \tag{2.2}$$

This suggests that for reversible walks if α is small then the Poincaré bound is often too pessimistic. We will see a similar phenomenon is true even for non-reversible walks.

When a general finite Markov has little or no holding probability then the $\frac{\rho_e}{\alpha}$ term will blow up. Then canonical paths can be used to bound the spectral gap λ_{PP^*} of PP^* by taking canonical paths along edges of walk PP^*, i.e. paths of even length with edges drawn alternately from P and then P^*. This leads to a bound of

$$\tau_x(\epsilon) \le 2\rho_e^{\mathsf{PP}^*} \min\{\rho_e^{\mathsf{PP}^*}, \ell\} \log\frac{1}{\epsilon\sqrt{\pi(x)}} \tag{2.3}$$

However, the minimal transition probability of PP^* is usually the square of that of P (e.g. usually $\min \mathsf{PP}^*(x,y) = \min \mathsf{P}(x,y)^2$ if π is uniform), making $\rho_e^{\mathsf{PP}^*}$ quite large.

3 Mixing bounds with r-conductance

Our extensions of past conductance results will be in terms of a measure which counts ergodic flow only up to some ratio r of the vertex size. When ergodic flow is well distributed among vertices then it may be possible to make r quite small without decreasing the ergodic flow significantly, and so the choice of r will measure some form of vertex congestion.

Definition 3.1. If $A, B \subset V$ and $r \in [0, 1/2]$ then the *r-ergodic flow* from A to B is

$$\mathsf{Q}_r(A, B) = \sum_{y \in B} \pi(y) \min \left\{ \frac{\mathsf{Q}(A, y)}{\pi(y)}, r \right\} .$$

The *r-ergodic flow* of A is $\mathsf{Q}_r(A) = \min\{\mathsf{Q}_r(A, A^c), \mathsf{Q}_r(A^c, A)\}$. The *$r$-conductance profile* is

$$\tilde{\Phi}_r(s) = \inf_{\pi(A) \in (0,s]} \tilde{\Phi}_r(A) \qquad \text{where} \qquad \tilde{\Phi}_r(A) = \frac{\mathsf{Q}_r(A)}{\pi(A)\pi(A^c)}$$

if $s \in (0, 1/2]$, while if $s > 1/2$ then $\tilde{\Phi}_r(s) = \tilde{\Phi}_r(1/2)$.

When $r \geq 1 - \alpha$ then $\mathsf{Q}_r(A) = \mathsf{Q}(A, A^c)$, and so for a lazy walk $\mathsf{Q}_{1/2}(A) = \mathsf{Q}(A, A^c)$ and $\tilde{\Phi}_{1/2}(A) = \tilde{\Phi}(A)$ agree with standard notions of ergodic flow and conductance. More generally, if ergodic flow from A to A^c and from A^c to A are not unduly concentrated on a few vertices then $\tilde{\Phi}_r(A) \approx \tilde{\Phi}(A)$ even for fairly small r. In such a situation our first result can lead to a substantially improved mixing bound:

Theorem 3.2. *For a finite ergodic Markov chain if $r \leq \alpha$ then*

$$\tau_x(\epsilon) \leq \left\lceil \int_{4\pi(x)}^{4/\epsilon^2} \frac{4 \min\{r, 1-r\}\, ds}{s \tilde{\Phi}_r(s)^2} \right\rceil .$$

This can be a factor $1/\alpha r$ times better than the Spectral profile result, although that used normal conductance instead of r-conductance. As will be discussed later, a natural application is a canonical path argument in which the paths are well distributed over the vertices.

Our second result is similar to the first, but for walks with no holding probability. As such we require an expression resembling modified conductance, but also having an r cutoff.

Definition 3.3. If $A \subset V$ and $r \in [0, 1]$ the *r-modified ergodic flow* is

$$\Psi_r(A) = \mathsf{Q}_r(A, V) - r\pi(A) = \sum_{y \in V} \pi(y) \min \left\{ \frac{\mathsf{Q}(A, y)}{\pi(y)}, r \right\} - r\pi(A) .$$

The *r-modified conductance profile* is given by

$$\tilde{\phi}^r(s) = \min_{\pi(A) \in (0,s]} \tilde{\phi}^r(A) \qquad \text{where} \qquad \tilde{\phi}^r(A) = \frac{\max\{\Psi_r(A), \Psi_r(A^c)\}}{\pi(A)\pi(A^c)} ,$$

if $s \in (0, 1/2]$, while if $s > 1/2$ then $\tilde{\phi}^r(s) = \tilde{\phi}^r(1/2)$.

To gain some intuition into this definition, note that if $\pi(A_r) \geq \pi(A)$ (see Definition 3.5) then $\Psi_r(A)$ is the minimal r-ergodic flow from A into a set of size $\pi(A^c)$. In particular, if $r \leq \alpha$ then $\Psi_r(A) = \mathsf{Q}_r(A, A^c)$. More generally, $\Psi_r(A) \leq \mathsf{Q}_r(A, A^c)$ always.

This time we can show a mixing time bound requiring no holding probability:

Theorem 3.4. *For a finite ergodic Markov chain if* $r \in (0,1]$ *then*

$$\tau_x(\epsilon) \leq \left\lceil \int_{4\pi(x)}^{4/\epsilon^2} \frac{12r\,ds}{s\min\{\tilde{\phi}^r(s)^2, r\tilde{\phi}^r(s)\}} \right\rceil$$

This improves on Theorem 3.2 up to a constant factor because the bound is decreasing in $\tilde{\phi}^r(s)$, and so if $r \leq \min\{\alpha, 1/2\}$ then $\tilde{\phi}^r(s) = \tilde{\Phi}_r(s) \leq 2r$. The case of $1/2 < r \leq \alpha$ reduces to the case of $1-r$, since $\min\{r, 1-r\} \geq 1-\alpha$ and so $\mathsf{Q}_{1-r}(A) = \mathsf{Q}(A, A^c) = \mathsf{Q}_r(A)$.

The unusual minimum in the denominator is necessary. With only the $\tilde{\phi}^r(s)^2$ term the lazy simple random walk on a complete graph K_n with loops (i.e. $\mathsf{P}(i,j) = \frac{1}{2n}$ if $i \neq j$ and $\mathsf{P}(i,i) = \frac{1}{2} + \frac{1}{2n}$) at $r = \min\{s/2, 1/2\}$ (the proof allows r to be increasing in s) would give the incorrect $\tau_x(1/2) = O(1)$, whereas with only the $r\tilde{\phi}^r(s)$ term the simple random walk $\mathsf{P}(i, i \pm 1) = 1/2$ on a cycle $\mathbb{Z}_n$ for n odd at $r = 1/2$ would give the incorrect $\tau_x(1/2) = O(n)$.

The proofs of both new bounds on mixing times are shown by use of the Evolving set methodology [9, 8]. To understand the method requires some new terminology:

Definition 3.5. Given $A \subset V$ and $u \in [0,1]$ let

$$A_u = \{y \in V \mid \mathsf{Q}(A,y) \geq u\pi(y)\} = \{y \in V \mid \mathsf{P}^*(y,A) \geq u\}\,.$$

The *root profile* $\tilde{\psi} : [0,1] \to [0,1]$ is any function satisfying

$$\tilde{\psi}(s) \leq \min_{\pi(A) \leq s} \tilde{\psi}(A) \quad \text{where} \quad \tilde{\psi}(A) = 1 - \frac{\int_0^1 \sqrt{\pi(A_u)(1-\pi(A_u))}\,du}{\sqrt{\pi(A)(1-\pi(A))}}$$

when $s \in (0, 1/2]$ and $\tilde{\psi}(s) \leq \tilde{\psi}(1/2)$ when $s \geq 1/2$.

Since $\mathsf{Q}(A, v) = \pi(v)$ then the set A_u indicates the set of vertices receiving at u-fraction of their inward flow from A. If $\pi(A_u)$ differs significantly from $\pi(A)$ when $u \in_{uar} [0,1]$ then this indicates the random walk expands quickly from set A into the full space. Since $\int_0^1 \pi(A_u)\,du = \sum_{y \in V} \pi(y) \frac{\mathsf{Q}(A,y)}{\pi(y)} = \mathsf{Q}(A, V) = \pi(A)$ and $f(a) = \sqrt{a(1-a)}$ is concave then $\int_0^1 f(\pi(A_u))\,du \leq f(\int_0^1 \pi(A_u)\,du) = f(\pi(A))$, and so large root profile indicates that $\pi(A_u)$ differs significantly from $\pi(A)$ if $u \in_{uar} [0,1]$, i.e. the random walk diffuses quickly. Morris and Peres [9] related $\psi(A) = 1 - \frac{\int_0^1 \sqrt{\pi(A_u)}\,du}{\sqrt{\pi(A)}}$ to mixing time, although we use a sharper form of [8] (found after equation (4.5)):

Theorem 3.6. *For a finite ergodic Markov chain*

$$\tau_x(\epsilon) \leq \begin{cases} \left\lceil \int_{4\pi(x)}^{4/\epsilon^2} \frac{ds}{s\,\tilde{\psi}(s)} \right\rceil & \textit{in general} \\ \left\lceil \int_{\pi(x)}^{1/\epsilon^2} \frac{ds}{2s\tilde{\psi}(s)} \right\rceil & \textit{if } s\tilde{\psi}\left(\frac{1}{1+s^2}\right) \textit{ is convex} \end{cases}$$

Evaluating $\tilde{\psi}(A)$ is primarily an optimization problem subject to whatever assumptions are given in the problem to be solved. It is often useful to recall that $\int_0^1 \pi(A_u)\,du = \pi(A)$ and note that $\pi(A_u)$ is a decreasing function of u.

To show Theorem 3.2 it then suffices to lower bound $\tilde{\psi}(A)$ when $r \leq \alpha$. Actually, when $1/2 \leq r \leq \alpha$ then $\min\{r, 1-r\} \geq 1-\alpha$ and so $\mathsf{Q}_{1-r}(A) = \mathsf{Q}(A, A^c) = \mathsf{Q}_r(A)$. Hence it suffices to consider the case when $r \leq \min\{\alpha, 1/2\}$.

Lemma 3.7. *If $A \subset V$ and $r \leq \min\{\alpha, 1/2\}$ then*

$$\tilde{\psi}(A) \geq 2r\left(1 - \sqrt{1 - \tilde{\Phi}_r(A)^2/4r^2}\right) \geq \frac{\tilde{\Phi}_r(A)^2}{4r}\,.$$

Proof. First consider the conditions of the lemma. If $v \in A$ then $\mathsf{Q}(A, v) \geq \mathsf{Q}(v, v) \geq \alpha\pi(v)$, and so if $u < r \leq \alpha$ then $A_u = A \cup \{y \in A^c \mid \mathsf{Q}(A, y) \geq u\pi(y)\}$. Hence

$$\begin{aligned}\int_0^r \pi(A_u)\,du &= r\pi(A) + \sum_{y \in A^c} \pi(y) \min\left\{\frac{\mathsf{Q}(A,y)}{\pi(y)},\, r\right\} \\ &= r\pi(A) + \mathsf{Q}_r(A, A^c) \geq r\pi(A) + \mathsf{Q}_r(A)\,.\end{aligned}$$

Likewise, $\int_{1-r}^1 \pi(A_u)\,du \leq r\pi(A) - \mathsf{Q}_r(A)$.

Now for the optimization. Let $x = \int_0^r (\pi(A_u) - \pi(A))\,du$ and $y = \int_{1-r}^1 (\pi(A) - \pi(A_u))\,du$. Note that $x, y \geq \mathsf{Q}_r(A)$, while $\int_r^{1-r} (\pi(A_u) - \pi(A))\,du = y - x$ because $\int_0^1 \pi(A_u)\,du = \pi(A)$. By Jensen's Inequality $\int f(g)\,d\mu \leq f(\int g\,d\mu)$ if f is concave and μ is a probability measure, and so if $f(a) = \sqrt{a(1-a)}$ then

$$\begin{aligned}\int_0^1 f(\pi(A_u))\,du &= r\int_0^r f(\pi(A_u))\,\frac{du}{r} + r\int_{1-r}^1 f(\pi(A_u))\,\frac{du}{r} + (1-2r)\int_r^{1-r} f(\pi(A_u))\,\frac{du}{1-2r} \\ &\leq r\,f(\pi(A) + x/r) + r\,f(\pi(A) - y/r) + (1-2r)\,f\left(\pi(A) + \frac{y-x}{1-2r}\right) =: h(x,y)\end{aligned}$$

By concavity f' is non-increasing, and so $\frac{\partial h}{\partial y}\Big|_{y \geq x} = -f'(\pi(A) - \frac{y}{r}) + f'(\pi(A) + \frac{y-x}{1-2r}) \leq 0$. Likewise, $\frac{\partial h}{\partial x}\Big|_{y \leq x} = f'(\pi(A) + \frac{x}{r}) - f'(\pi(A) + \frac{y-x}{1-2r}) \leq 0$. Hence h is maximized when $y = x$. But $\frac{d}{dx}h(x,x) = f'(\pi(A) + \frac{x}{r}) - f'(\pi(A) - \frac{x}{r}) \leq 0$, and so $\int_0^1 f(\pi(A_u))\,du \leq h(x,y) \leq h(\mathsf{Q}_r(A), \mathsf{Q}_r(A))$.

Putting together these various facts we have that

$$\begin{aligned}1 - \tilde{\psi}(A) &\leq 1 - 2r + r\frac{f(\pi(A) + \mathsf{Q}_r(A)/r) + f(\pi(A) - \mathsf{Q}_r(A)/r)}{f(\pi(A))} \\ &= 1 - 2r + 2r\left(\sqrt{\left(\frac{1}{2} + \frac{\mathsf{Q}_r(A)}{2r\pi(A)}\right)\left(\frac{1}{2} - \frac{\mathsf{Q}_r(A)}{2r\pi(A^c)}\right)} + \sqrt{\left(\frac{1}{2} - \frac{\mathsf{Q}_r(A)}{2r\pi(A)}\right)\left(\frac{1}{2} + \frac{\mathsf{Q}_r(A)}{2r\pi(A^c)}\right)}\right) \\ &\leq 1 - 2r + 2r\sqrt{1 - \tilde{\Phi}_r(A)^2/4r^2} \leq 1 - \tilde{\Phi}_r(A)^2/4r\end{aligned}$$

by Lemma 3.8 with $X = \frac{1}{2} + \frac{\mathsf{Q}_r(A)}{2r\pi(A)}$ and $Y = \frac{1}{2} - \frac{\mathsf{Q}_r(A)}{2r\pi(A^c)}$, and lastly $\sqrt{1-x} \leq 1 - x/2$. $\square$

The following small inequality was required in the proof.

Lemma 3.8. *If* $X,\, Y \in [0,1]$ *then*

$$g(X,Y) = \sqrt{X\,Y} + \sqrt{(1-X)(1-Y)} \leq \sqrt{1-(X-Y)^2}\,.$$

Proof. Observe that

$$g(X,Y)^2 = 1-(X+Y)+2\,X\,Y + \sqrt{[1-(X+Y)+2\,X\,Y]^2 - [1-2(X+Y)+(X+Y)^2]}\,.$$

Now, $\sqrt{A^2-B} \leq A-B$ if $A^2 \geq B$, $A \leq \frac{1+B}{2}$ and $A \geq B$ (square both sides to show this). These conditions are easily verified with $A = 1-(X+Y)+2\,X\,Y$ and $B = 1-2(X+Y)+(X+Y)^2$, and so

$$\begin{aligned} g(X,Y)^2 &\leq 2\,[1-(X+Y)+2\,X\,Y] - \left[1-2(X+Y)+(X+Y)^2\right] \\ &= 1+2\,X\,Y - X^2 - Y^2 = 1-(X-Y)^2 \end{aligned}$$

□

Finally, the following lemma suffices to prove Theorem 3.4:

Lemma 3.9. *If* $A \subset V$ *and* $r \in (0,1]$ *then*

$$\tilde{\psi}(A) \geq \frac{1}{12r}\,\min\left\{\tilde{\phi}^r(A)^2,\ r\tilde{\phi}^r(A)\right\}$$

Proof. It suffices to assume that $r \in (0,1)$ because the case of $r=1$ follows by taking a limit.

Observe that $(A_u)^c = (A^c)_{1-u}$ for every u such that $\forall v \in V:\ \mathsf{Q}(A,v) \neq u\pi(v)$, and in particular u-a.e., so $\tilde{\psi}(A) = \tilde{\psi}(A^c)$. Moreover $\tilde{\phi}^r(A) = \tilde{\phi}^r(A^c)$. The lemma then holds for set A if and only if it holds for A^c. Thus, without loss we assume that $\Psi_r(A) \geq \Psi_r(A^c)$, since if $\Psi_r(A^c) > \Psi_r(A)$ then $\Psi_r(B) > \Psi_r(B^c)$ for $B = A^c$.

As before, the quantities in the lemma can be rewritten in terms of A_u:

$$\int_0^r \pi(A_u)\,du = \sum_{y\in V} \pi(y)\min\left\{\frac{\mathsf{Q}(A,y)}{\pi(y)},\ r\right\} = r\pi(A) + \Psi_r(A)$$

Since $\int_0^1 \pi(A_u)\,du = \pi(A)$ then $\int_r^1 \pi(A_u)\,du = (1-r)\pi(A) - \Psi_r(A)$.

Consider the case when $\tilde{\phi}^r(A) = \frac{\Psi_r(A)}{\pi(A)\pi(A^c)} \leq 1-2r$. Then

$$\begin{aligned} 1-\tilde{\psi}(A) &= \frac{r\int_0^r \sqrt{\pi(A_u)(1-\pi(A_u))}\,\frac{du}{r} + (1-r)\int_r^1 \sqrt{\pi(A_u)(1-\pi(A_u))}\,\frac{du}{1-r}}{\sqrt{\pi(A)(1-\pi(A))}} \\ &\leq r\sqrt{\left(1+\frac{\tilde{\phi}^r(A)(1-\pi(A))}{r}\right)\left(1-\frac{\tilde{\phi}^r(A)\pi(A)}{r}\right)} \\ &\quad +(1-r)\sqrt{\left(1-\frac{\tilde{\phi}^r(A)(1-\pi(A))}{1-r}\right)\left(1+\frac{\tilde{\phi}^r(A)\pi(A)}{1-r}\right)} \\ &\leq r\sqrt{1+\frac{\tilde{\phi}^r(A)}{r}} + (1-r)\sqrt{1-\frac{\tilde{\phi}^r(A)}{1-r}} \\ &\leq r\left(1+\frac{\tilde{\phi}^r(A)}{2r} - \frac{1}{12}\min\left\{\frac{\tilde{\phi}^r(A)^2}{r^2},\,\frac{\tilde{\phi}^r(A)}{r}\right\}\right) + (1-r)\left(1-\frac{\tilde{\phi}^r(A)}{2(1-r)}\right) \end{aligned}$$

The first inequality is Jensen's with $f(x) = \sqrt{x(1-x)}$ and some simplification, the second is because if $\tilde{\phi}^r(A) \leq 1-2r$ then the equation is concave as a function of $x = \pi(A) \in [0, 1/2]$ with negative derivative at $x = 0$ and hence maximum at $x = 0$, while the final line applies the relations $\sqrt{1+x} \leq 1 + \frac{x}{2} - \frac{\min\{x^2,x\}}{12}$ and $\sqrt{1-x} \leq 1 - \frac{x}{2}$ if $x \geq 0$. This simplifies to give the lemma.

When $\tilde{\phi}^r(A) > 1-2r$ then the maximum in the second inequality above is instead at $x = \pi(A) = \frac{1-(1-2r)/\tilde{\phi}^r(A)}{2}$, and so

$$1 - \tilde{\psi}(A) \leq \sqrt{1 - \tilde{\phi}^r(A)^2} \leq 1 - \frac{\tilde{\phi}^r(A)^2}{2}\,.$$

If $r \geq 1/3$ then $\frac{\tilde{\phi}^r(A)^2}{2} \geq \frac{\tilde{\phi}^r(A)^2}{6r}$, while if $r < 1/3$ then $\frac{\tilde{\phi}^r(A)^2}{2} > \frac{(1-2r)\tilde{\phi}^r(A)}{2} > \frac{r\tilde{\phi}^r(A)}{6r}$. In either case the above equation establishes the lemma. □

4 New Canonical path bounds

When ergodic flow leaving a set is not heavily concentrated at a few vertices then our r-conductance can improve substantially over the conductance method of [5]. In this section we explore such a situation by considering the case when canonical paths are well distributed among the vertices, and thus the ergodic flow induced by canonical paths is not heavily concentrated at any vertex. A notion of vertex congestion will thus be required:

Definition 4.1. Given paths Γ the *vertex congestion* is

$$\rho_v = \rho_v(\Gamma) = \max_{v \in V} \frac{1}{\pi(v)} \left(\sum_{\substack{x,y: \\ v \in \gamma_{xy} \setminus \gamma_{yx}}} \pi(x)\pi(y) + \frac{1}{2} \sum_{\substack{x,y: \\ v \in \gamma_{xy} \cap \gamma_{yx}}} \pi(x)\pi(y) \right).$$

The *minimum boundary probability* is

$$\mathsf{P}_0(\Gamma) = \min\{\mathsf{P}(a,b) \mid \exists x,y : (a,b) \in \gamma_{xy}\}\,.$$

This form was chosen to ensure that a vertex v is only counted in one of the paths γ_{xy} or γ_{yx}, and so undirected paths are counted only once. To simplify this note that $v \in \gamma_{vx} \cap \gamma_{xv}$ for all $x \in V$, and so

$$\rho_v \leq \max_{v \in V} \frac{1}{\pi(v)} \sum_{\gamma_{xy} \ni v \neq x} \pi(x)\pi(y)\,. \tag{4.4}$$

Also, in an undirected graph if γ_{yx} is just γ_{xy} traversed in reverse then

$$\rho_v = \frac{1}{2} \max_{v \in V} \frac{1}{\pi(v)} \sum_{\gamma_{xy} \ni v} \pi(x)\pi(y)\,. \tag{4.5}$$

It will be seen that $\tilde{\Phi}_{\rho_v/\rho_e} \geq 1/\rho_e$, and from this we derive our first path bound:

Theorem 4.2. *For a finite ergodic Markov chain, if* $\pi_0 = \min \pi(v)$ *and* $\epsilon \leq \sqrt{2}$ *then*

$$\begin{aligned} \tau_x(\epsilon) &\leq 4\rho_v \max\left\{\frac{\rho_v}{\alpha},\, \rho_e\right\} \log \frac{1}{\epsilon\sqrt{\pi(x)}} \\ \tau_x(\epsilon) &\leq \frac{\rho_v^2 - 1}{\min\{\alpha, \mathsf{P}_0(\Gamma)\}} + 4\rho_v \max\left\{\frac{\rho_v}{\alpha},\, \rho_e\right\} \log \frac{1}{\epsilon\sqrt{\pi_0 \rho_v}} \end{aligned}$$

Remark 4.3. *Since* $\rho_v \leq \rho_e(1-\alpha)$ *and* $\rho_e \leq \frac{\rho_v}{\mathsf{P}_0(\Gamma)}$ *then* $\max\left\{\frac{\rho_v}{\alpha}, \rho_e\right\} \leq \frac{1-\alpha}{\alpha}\,\rho_e,\ \frac{\rho_v}{\min\{\alpha,\mathsf{P}_0(\Gamma)\}}$.

The maximum and $\pi_0\rho_v$ terms are not simply artifacts of the proof. A simple illustration of this is the walk on a cycle $\mathbb{Z}_n$ with $\mathsf{P}(i,i) = \alpha \in (0,1)$ and $\mathsf{P}(i,i+1) = 1-\alpha$. With the obvious choice of paths this has $\rho_v = \frac{n-1}{2}$, $\rho_e = \frac{n-1}{2(1-\alpha)}$, $\mathsf{P}_0(\Gamma) = 1-\alpha$, $\pi_0 = 1/n$ and mixing time $\tau_x(\epsilon) = \Theta\left(\frac{n^2}{\min\{\alpha,1-\alpha\}} \log \frac{1}{\epsilon}\right)$.

When there is no holding probability then we can show a bound which combines aspects of both the non-lazy canonical path methods discussed in the preliminaries, in that we require paths of odd length with edges drawn alternately from P and then P^*. A slightly different vertex congestion is required in this case:

Definition 4.4. For each $x, y \in V$ (including $x = y$) define an odd length path $\dot\gamma_{xy}$ from x to y alternating between edges of P and P^*,

$$x = x_0 \xrightarrow{\mathsf{P}} x_1 \xrightarrow{\mathsf{P}^*} x_2 \xrightarrow{\mathsf{P}} x_3 \xrightarrow{\mathsf{P}^*} \cdots \xrightarrow{\mathsf{P}} x_{2n+1} = y\,.$$

Denote $(u,v) \in_{\mathsf{P}} \gamma_{xy}$ if $u \xrightarrow{\mathsf{P}} v$ or $v \xrightarrow{\mathsf{P}^*} u$ appear in path $\dot\gamma_{xy}$, i.e. the directed edge (u,v) of P was used in constructing the path (as edges of P^* are reversals of edges of P). Also, $v \in_{odd} \dot\gamma_{xy}$ if some $(u,v) \in_{\mathsf{P}} \dot\gamma_{xy}$, i.e. $v = x_{2i+1}$ for some $i \in [0,n]$.

Let $\dot\Gamma$ be the collection of these paths. The *vertex-congestion* is

$$\dot\rho_v = \dot\rho_v(\dot\Gamma) = \max_{v\in V} \frac{1}{\pi(v)} \sum_{\{x,y \,|\, v \in_{odd} \dot\gamma_{xy}\}} \pi(x)\pi(y)\,.$$

The *minimum reversed boundary probability* is

$$\mathsf{P}_0^*(\dot\Gamma) = \min\{\mathsf{P}^*(b,a) \mid \exists x,y : (a,b) \in_{\mathsf{P}} \dot\gamma_{xy}\}\,.$$

We will find that $\tilde\phi^{\mathsf{P}_0^*(\dot\Gamma)} \geq \frac{\mathsf{P}_0^*(\dot\Gamma)}{2\dot\rho_v}$, from which we derive our general result:

Theorem 4.5. *For a finite ergodic Markov chain, if* $\epsilon \leq 1$ *then*

$$\tau_x(\epsilon) \leq \left\lceil \frac{40\dot\rho_v^2}{\mathsf{P}_0^*(\dot\Gamma)} \log \frac{1}{\epsilon\sqrt{\delta_0\dot\rho_v}} \right\rceil$$

where $\delta_0 = \min_{A,B\subset V}\{|\pi(A) - \pi(B)| \mid \pi(A) \neq \pi(B)\}$.

The requirement of odd length paths is necessary because a simple random walk on a cycle with an even number of vertices does not converge, while paths must alternate between P and P^* because a clockwise walk on a cycle of odd length (i.e. $\mathsf{P}(i,i+1) = 1$) does not converge.

This is not directly comparable to Theorem 4.2. However, given paths Γ along edges of P one can define paths $\dot\gamma_{xy} \in \dot\Gamma$ by taking the P steps to be self-loops and the P^* steps to be the reverse of path γ_{yx} (define $\gamma_{xx} = \emptyset$ so $\dot\gamma_{xx}$ is a self-loop). Then

$$\dot\rho_v = 1 + \max_{v\in V} \frac{1}{\pi(v)} \sum_{\gamma_{xy}\ni v\neq x} \pi(x)\pi(y)$$

while $\mathsf{P}_0^*(\dot\Gamma) = \min\{\alpha, \mathsf{P}_0(\Gamma)\}$, and so Theorem 4.5 resembles the second bound of Theorem 4.2 under the simplifications suggested in (4.4) and Remark 4.3.

Remark 4.6. *It may seem odd that for a lazy walk there are three forms of the mixing bound: the original conductance bound involves $\rho_e\rho_e$, the Poincaré bound uses $\rho_e\ell$, while our bound has $\rho_e\rho_v$. Certainly $\rho_v \leq \rho_e$ so our bound strictly improves on the original conductance result. Also, if we define average vertex congestion and average path length by*

$$\rho_v^{ave} = \sum_{v\in V} \pi(v) \left[\frac{1}{\pi(v)} \sum_{\gamma_{xy}\ni v\neq x} \pi(x)\,\pi(y)\right] \qquad \text{and} \qquad \ell_{ave} = \frac{\sum_{x\neq y\in V} \pi(x)\,\pi(y)\,|\gamma_{xy}|}{\sum_{x\neq y\in V} \pi(x)\pi(y)}$$

respectively, then

$$\begin{aligned} \rho_v^{ave} = & \sum_{x\neq y\in V} \sum_{\gamma_{xy}\ni v\neq x} \pi(x)\,\pi(y) &= \sum_{x\neq y\in V} \pi(x)\,\pi(y)\,|\gamma_{xy}| \\ = & \quad \ell_{ave} \sum_{x\neq y\in V} \pi(x)\,\pi(y) &= \ell_{ave}(1 - \|\pi\|_2^2)\ . \end{aligned}$$

Hence vertex-congestion and path length are of similar sizes when the canonical paths are well distributed among the vertices. When a few paths are very long, or the distribution is concentrated at a single vertex, then vertex-congestion can be much smaller than maximum path length, although it is more often the case for maximum path length to be smaller than vertex-congestion.

To prove Theorem 4.2 we require a lower bound on an appropriate choice of r-conductance.

Lemma 4.7. *Given cycle-free paths $\{\gamma_{xy}\}$ between each $x, y \in V$, $x \neq y$, then*

$$\tilde\Phi_{\rho_v/\rho_e}(A) \geq 1/\rho_e\,.$$

Proof. For each pair of vertices $x \in A$, $y \in A^c$, $x \neq y$, transport flow of $\pi(x)\pi(y)$ along the path γ_{xy} from x to y, for a total of $\pi(A)\pi(A^c)$ from A to A^c. Then, if $b \in A^c$,

$$\pi(b) \geq \frac{1}{\rho_v} \sum_{\substack{(x,y)\in A\times A^c:\\ b\in\gamma_{xy}}} \pi(x)\pi(y) = \frac{1}{\rho_v} \sum_{a\in A} \sum_{\substack{(x,y)\in A\times A^c:\\ (a,b)\in\gamma_{xy}}} \pi(x)\pi(y)\,.$$

Likewise, if $a \in A$ and $b \in A^c$ then

$$\pi(a)\mathsf{P}(a,b) \geq \frac{1}{\rho_e} \sum_{\substack{(x,y)\in A\times A^c:\\ (a,b)\in\gamma_{xy}}} \pi(x)\pi(y)\,.$$

Hence,

$$\begin{aligned}
\mathsf{Q}_{\rho_v/\rho_e}(A, A^c) &= \sum_{b\in A^c} \min\left\{\sum_{a\in A} \pi(a)\mathsf{P}(a,b),\ \frac{\rho_v}{\rho_e}\pi(b)\right\} \\
&\geq \sum_{b\in A^c}\sum_{a\in A}\frac{1}{\rho_e}\sum_{\substack{(x,y)\in A\times A^c:\\ (a,b)\in\gamma_{xy}}} \pi(x)\pi(y) \\
&= \frac{1}{\rho_e}\sum_{(x,y)\in A\times A^c} \pi(x)\pi(y)\ |\{(a,b)\in A\times A^c \mid (a,b)\in\gamma_{xy}\}| \\
&\geq \frac{1}{\rho_e}\sum_{(x,y)\in A\times A^c} \pi(x)\pi(y) = \frac{\pi(A)\pi(A^c)}{\rho_e}
\end{aligned}$$

A similar argument lower bounds $\mathsf{Q}_{\rho_v/\rho_e}(A^c, A)$. □

It is then not hard to show our first canonical path result:

Proof of Theorem 4.2. Remove any cycles in the paths. This can only decrease ρ_v and ρ_e, and hence improve the mixing time bound. If $\rho_v/\rho_e \leq \alpha$ then $\tilde{\psi}(A) \geq 1/4\rho_v\rho_e$ by Lemmas 3.7 and 4.7. If, instead, $\rho_v/\rho_e > \alpha$ then at least $\tilde{\Phi}_\alpha(A) \geq \frac{\alpha}{\rho_v/\rho_e}\tilde{\Phi}_{\rho_v/\rho_e}(A) \geq \alpha/\rho_v$, and so $\tilde{\psi}(A) \geq \alpha/4\rho_v^2$. Hence

$$\tilde{\psi}(A) \geq \frac{1}{4\rho_v \max\{\rho_v/\alpha,\ \rho_e\}}\,.$$

The convexity condition of Theorem 3.6 is satisfied by this lower bound, and so the first mixing time bound follows by a simple integration.

For the second mixing time bound we improve bounds on r-ergodic flow from small sets. Since $\mathsf{Q}_{\rho_v/\rho_e}(A, A^c) > 0$ was shown using only edges in paths of Γ, there exists $(a,b), (x,y) \in A\times A^c$ with $(a,b)\in\gamma_{xy}$, and so $\mathsf{Q}(A,b) \geq \mathsf{Q}(a,b) \geq \pi_0\mathsf{P}_0(\Gamma)$. Then $\mathsf{Q}_{\mathsf{P}_0(\Gamma)}(A) \geq \pi_0\mathsf{P}_0(\Gamma)$, and similarly $\mathsf{Q}_{\mathsf{P}_0(\Gamma)}(A^c) \geq \pi_0\mathsf{P}_0(\Gamma)$. It follows that $\tilde{\Phi}_{\mathsf{P}_0(\Gamma)} \geq \frac{\pi_0\mathsf{P}_0(\Gamma)}{\pi(A)\pi(A^c)}$, and as in the previous paragraph then $\tilde{\Phi}_{\min\{\alpha,\mathsf{P}_0(\Gamma)\}} \geq \frac{\pi_0\min\{\alpha,\mathsf{P}_0(\Gamma)\}}{\pi(A)\pi(A^c)}$. By Lemma 3.7,

$$\tilde{\psi}(A) \geq \max\left\{\frac{1}{4\rho_v\max\{\rho_v/\alpha,\ \rho_e\}},\ \frac{\min\{\alpha,\mathsf{P}_0(\Gamma)\}\pi_0^2}{4\pi(A)^2\pi(A^c)^2}\right\}\,.$$

The convexity condition of Theorem 3.6 is easily verified for this lower bound and so

$$\tau_x(\epsilon) \leq \int_{\pi_0}^{\pi_0\rho_v} \frac{2s\,ds}{\min\{\alpha,\mathsf{P}_0(\Gamma)\}\,\pi_0^2} + 2\rho_v\max\left\{\frac{\rho_v}{\alpha},\ \rho_e\right\}\int_{\pi_0\rho_v}^{1/\epsilon^2}\frac{ds}{s}$$

if $\pi_0\rho_v \leq 1/\epsilon^2$, which integrates to give the second bound in the theorem. Since $\pi_0\rho_v \leq \pi_0\frac{1/2}{\pi_0} = \frac{1}{2}$ then it suffices that $\epsilon \leq \sqrt{2}$. □

For the canonical path theorem without holding probabilities we require a lower bound on an appropriate choice of r-modified conductance.

Lemma 4.8. *If A is a proper subset of V then*

$$\tilde{\phi}^{\mathsf{P}_0^*(\dot{\Gamma})}(A) \geq \frac{\mathsf{P}_0^*(\dot{\Gamma})}{2\dot{\rho}_v}\,.$$

Proof. Define $E = \{(x,y) \mid \exists a, b : (x,y) \in_{\mathsf{P}} \dot{\gamma}_{ab}\}$ to be the edge set used in paths. Since $v \in_{odd} \dot{\gamma}_{xv}$ then every vertex is an endpoint in E and we may partition the state space into vertices with incoming edges in E from A, A^c, or both: $B = \{v \in V \mid (x,v) \in E \Rightarrow x \in A\}$, $C = \{v \in V \mid (x,v) \in E \Rightarrow x \in A^c\}$, and $D = V \setminus (B \cup C) = \{v \in V \mid \exists x \in A,\, y \in A^c;\, (x,v) \in E;\, (y,v) \in E\}$. Then

$$\Psi_{\mathsf{P}_0^*(\dot{\Gamma})}(A) \geq \mathsf{Q}_{\mathsf{P}_0^*(\dot{\Gamma})}(A, B \cup D) - \mathsf{P}_0^*(\dot{\Gamma})\pi(A) = \mathsf{P}_0^*(\dot{\Gamma})(\pi(B \cup D) - \pi(A))\,.$$

Likewise, $\Psi_{\mathsf{P}_0^*(\dot{\Gamma})}(A^c) \geq \mathsf{P}_0^*(\dot{\Gamma})(\pi(C \cup D) - \pi(A^c)) = \mathsf{P}_0^*(\dot{\Gamma})(\pi(A) - \pi(B))$.

We now lower bound $\pi(D)$. If $x \in A$ and $y \in V \setminus B$ then on path $\dot{\gamma}_{xy}$ if $x_{2i+1} \in B$ then $x_{2i+2} \in A$, so the first vertex $x_{2i+1} \notin B$ will be in D, and in particular $|\{a \in D : a \in_{odd} \dot{\gamma}_{xy}\}| \geq 1$. But then

$$\begin{aligned}
\pi(D) &= \sum_{a \in D} \pi(a) \geq \sum_{a \in D} \sum_{\substack{(x,y) \in A \times (V \setminus B): \\ a \in_{odd} \dot{\gamma}_{xy}}} \frac{\pi(x)\pi(y)}{\dot{\rho}_v} \\
&= \sum_{(x,y) \in A \times (V \setminus B)} \frac{\pi(x)\pi(y)}{\dot{\rho}_v} \left|\{a \in D \,:\, a \in_{odd} \dot{\gamma}_{xy}\}\right| \\
&\geq \frac{\pi(A)\pi(V \setminus B)}{\dot{\rho}_v}\,.
\end{aligned}$$

If $\pi(B) \leq \pi(A) - \frac{\pi(A)\pi(A^c)}{2\dot{\rho}_v}$ then $\Psi_{\mathsf{P}_0^*(\dot{\Gamma})}(A^c) \geq \mathsf{P}_0^*(\dot{\Gamma})\frac{\pi(A)\pi(A^c)}{2\dot{\rho}_v}$. On the other hand, if $\pi(B) > \pi(A) - \frac{\pi(A)\pi(A^c)}{2\dot{\rho}_v}$ then $\pi(B \cup D) - \pi(A) \geq \pi(B) + \frac{\pi(A)(1-\pi(B))}{\dot{\rho}_v} - \pi(A) > \frac{\pi(A)\pi(A^c)}{2\dot{\rho}_v}$. It follows that $\tilde{\phi}^{\mathsf{P}_0^*(\dot{\Gamma})}(A) \geq \frac{\mathsf{P}_0^*(\dot{\Gamma})}{2\dot{\rho}_v}$. ☐

Proof of Theorem 4.5. Observe that $\dot{\rho}_v \geq 1$ because $v \in_{odd} \dot{\gamma}_{uv}$ for every $u, v \in V$. Then proceed as in the second bound of Theorem 4.2, but with $\tilde{\phi}^{\mathsf{P}_0^*(\dot{\Gamma})}(A) \geq \frac{\Psi_{\mathsf{P}_0^*(\dot{\Gamma})}(A)}{\pi(A)} \geq \frac{\mathsf{P}_0^*(\dot{\Gamma})\delta_0}{\pi(A)}$ used when $\pi(A) \in (\pi_0, 2\delta_0\dot{\rho}_v]$, and Lemmas 3.9 and 4.8 in place of Lemmas 3.7 and 4.7. This gives a coefficient of 48. To reduce this to 40 observe that in the proof of Lemma 3.9 we may use the stronger relation $\sqrt{1+x} \leq 1 + \frac{x}{2} - \frac{x^2}{10}$ when $x \in [0, 1/2]$ to show that $\tilde{\psi}(A) \geq \frac{\tilde{\phi}^r(A)^2}{10r}$ when $\tilde{\phi}^r(A) \leq \frac{r}{2}$. ☐

5 Examples

We present two examples where our canonical path theorems extend previously known bounds into the general non-reversible non-lazy case. First, the classical problem of the max-degree walk on a graph, then walks on Cayley graphs.

Example 5.1. Consider an Eulerian multigraph with n vertices and max-degree d, that is a strongly connected graph with in-degree=out-degree at each vertex. This is the natural generalization of an undirected multigraph into the directed graph setting. Let $d(x,y)$ denote the number of directed edges from x to y, so that $d(x) = \sum_y d(x,y)$ is the degree of x. The max-degree walk has uniform stationary distribution and transition probabilities $\mathsf{P}(x,y) = \frac{d(x,y)}{d}$ if $y \neq x$ and $\mathsf{P}(x,x) = 1 - \frac{d(x,y)-d(x,x)}{d}$; the walk is lazy if there are at least $d/2$ self-loops at each vertex.

By strong connectivity, for every $x \neq y \in V$ there is a shortest path γ_{xy} with $|\gamma_{xy}| \leq n-1$. In order to eliminate significant periodicity effects it is sufficient to assume that every vertex has a single self-loop, i.e. holding probability $\alpha \geq 1/d$. Then, even if every path passes through some vertex v and edge e, then $\rho_v \leq \frac{1}{2\pi(x)} = \frac{n}{2}$ (half size because $v \in \gamma_{xy} \cap \gamma_{yx}$) and $\rho_e \leq \frac{nd}{1-\alpha}$, and so by Theorem 4.2,

$$\tau_x(\epsilon) \leq \left\lceil dn^2 \log \frac{1}{\epsilon} \right\rceil .$$

In contrast, the Poincaré method (2.1) can show a comparable bound only for the lazy walk, while the odd path length (2.2) only for the reversible case (i.e. $d(x,y) = d(y,x)$ for every x, y). If (2.1) or (2.3) were used to study the Eulerian case then a factor of d would be lost in either case.

More generally, without the self-loop condition, if the graph is connected under paths $\dot{\gamma}_{xy}$ of odd length alternating between edges of $E = \{(x,y) : d(x,y) \geq 1\}$ and $E^* = \{(x,y) : d(y,x) \geq 1\}$ then $\dot{\rho}_v \leq \frac{1}{\pi(x)} = n$ and $\mathsf{P}_0^* \geq 1/d$. By Theorem 4.5,

$$\tau_x(\epsilon) \leq \left\lceil 40\, dn^2 \log \frac{1}{\epsilon} \right\rceil .$$

In comparison, (2.2) can show a comparable bound only in the reversible case, while if (2.3) is used to study this then at worst $\rho_e^{\mathsf{PP}^*} = \Theta(nd^2)$ and and so a factor of d is lost.

Our methods have thus shown that canonical path methods need not lose a factor of d when going from the undirected to the Eulerian directed case, both with the self-loop requirement and without.

A more interesting problem is walks on Cayley graphs, that is, random walks on groups.

Example 5.2. Consider a group G with (non-symmetric) generating set S. The Cayley graph has edge set (g, gs) for all $g \in G$, $s \in S$. If p is a probability distribution on S then consider the walk $\mathsf{P}(g, gs) = p(s)$. Babai [1] showed a mixing bound of $\tau_x(\epsilon) = O(\frac{\Delta^2}{\min_{s\in S} p(s)} \log(|G|/\epsilon))$ for the lazy walk on Cayley graphs (i.e. $p(id) \geq 1/2$) with symmetric generating set and diameter Δ. Diaconis and Saloff-Coste [2] used canonical paths to extend this to the non-symmetric lazy case, and to the symmetric case if merely $id \in S$. Our path methods show there is no need to differentiate between symmetric and non-symmetric cases, that is, the same bound holds for the non-symmetric case with $id \in S$; we show a similar bound holds even if $id \notin S$.

First the case of $id \in S$, so there is a holding probability. To each $g \in G$ write $g = s_1 s_2 \cdots s_k$ as a product of generators. Let $\Delta = \max |g|$ be the length of the longest such representation, and $N(g,s) \leq \Delta$ denote the number of times generator s appears in the representation of g. Then the argument of [2] is easily extended (see Appendix) to show that there are canonical paths such that

$$\rho_v < \Delta, \qquad \rho_e < \max_{g\in G,\, s\in S-id} \frac{N(g,s)}{p(s)} .$$

It follows from Theorem 4.2 that

$$\tau_x(\epsilon) \leq \left\lceil \frac{4\Delta^2}{\min_{s\in S} p(s)} \left(\frac{1}{2} \log \frac{2|G|}{\Delta} + \log \frac{1}{\epsilon} \right) \right\rceil .$$

This is (up to a small constant) the same bound as Babai, and Diaconis and Saloff-Coste showed in weaker settings. If (2.1) were used to extend their bounds to this more general setting it would have had the weaker $p(id)\min_{s\in S-id} p(s) \leq (\min_{s\in S} p(s))^2$ in the denominator.

Now suppose the identity is not in S. Then, if the graph has diameter Δ^* where diameter is measured along odd length paths alternating between elements of S and S^{-1}, the argument in the Appendix applies to $\dot{\rho}_v$ as well, but with $\frac{1+|\gamma_{xy}|}{2}$ vertices at odd positions of path γ_{xy}, and without $1-\pi(g)$ being subtracted. Hence $\dot{\rho}_v \leq \frac{1+\Delta^*}{2}$ and so

$$\tau_x(\epsilon) \leq \left\lceil \frac{10(1+\Delta^*)^2}{\min_{s\in S} p(s)} \left(\frac{1}{2}\log\frac{2|G|}{1+\Delta^*} + \log\frac{1}{\epsilon} \right) \right\rceil .$$

In comparison, (2.2) can show this only with symmetric generating set, while if (2.3) is used to study the general setting then it results in $\min_{ss'\in SS^*} p(s)p^*(s') = (\min_{s\in S} p(s))^2$ in the denominator.

Our methods have thus shown the squaring of $\min p(s)$ to be unnecessary when going from the symmetric to non-symmetric setting, both in the case of small holding probability and the case of no holding probability.

References

[1] L. Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing (STOC 1991)*, pages 164–174, 1991.

[2] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.

[3] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of markov chains. *The Annals of Applied Probability*, 1:36–61, 1991.

[4] S. Goel, R. Montenegro, and P. Tetali. Mixing time bounds via the spectral profile. *Electronic Journal of Probability*, 11:1–26, 2006.

[5] M. Jerrum and A. Sinclair. Conductance and the rapid mixing property for markov chains: the approximation of the permanent resolved. *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, pages 235–243, 1988.

[6] R. Kannan, L. Lovász, and R. Montenegro. Blocking conductance and mixing in random walks. *Combinatorics, Probability and Computing*, 15(4):541–570, 2006.

[7] L. Lovász and R. Kannan. Faster mixing via average conductance. *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC 1999)*, pages 282–287, 1999.

[8] R. Montenegro and P. Tetali. *Mathematical Aspects of Mixing Times in Markov Chains*, volume 1:3 of *Foundations and Trends in Theoretical Computer Science.* NOW Publishers, Boston-Delft, June 2006.

[9] B. Morris and Y. Peres. Evolving sets, mixing and heat kernel bounds. *Probability Theory and Related Fields*, 133(2):245–266, 2005.

[10] A. Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992.

Appendix

In order to study walks on Cayley graphs we claimed that a result of [2] generalizes easily. Details can be found below.

Lemma 5.3. *Consider group G with (non-symmetric) generating set S, and to each $g \in G$ write $g = s_1 s_2 \cdots s_k$ as a product of generators. Let $\Delta = \max|g|$ be the length of the longest such representation, and $N(g,s)$ denote the number of times generator s appears in the representation of g. If p is a probability distribution on S then the random walk with transitions $\mathsf{P}(g,gs) = p(s)$ has canonical paths such that*

$$\rho_v < \Delta, \qquad \rho_e < \max_{g\in G,\, s\in S} \frac{N(g,s)}{p(s)}\,.$$

If the generating set is symmetric then $\rho_v < \frac{\Delta+1}{2}$.

Proof. Given $x, y \in G$ let $g = x^{-1}y = s_1 s_2 \cdots s_k$ and path $\gamma_{x,y}$ be given by $x \to xs_1 \to \cdots \to xg = y$.

First vertex-congestion. The same number of paths pass through each vertex, because if $\gamma_{x,y}$ includes vertex v then $\gamma_{xv^{-1}w,yv^{-1}w}$ includes vertex w, and vice-versa. Then, by equation (4.4),

$$\begin{aligned}
\rho_v &\le \frac{1}{|G|}\sum_{g\in G}\left(\frac{1}{\pi(g)}\sum_{\substack{x,y\in G\,|\,x\neq y,\\ g\in\gamma_{xy}}}\pi(x)\pi(y) - (1-\pi(g))\right)\\
&\le \sum_{x,y\in G\,|\,x\neq y}(1+|\gamma_{xy}|)\,\pi(x)\pi(y) - \frac{|G|-1}{|G|}\\
&\le (1+\Delta)\left(1-\sum_{g\in G}\pi(g)^2\right) - 1 + \frac{1}{|G|} = \Delta - \frac{\Delta}{|G|}\,.
\end{aligned}$$

For the undirected case use the form of ρ_v given in equation (4.5).

Now edge-congestion. For fixed $s \in S$, the same number of paths pass through each edge (g, gs) $\forall g \in G$, because if $\gamma_{x,y}$ includes edge (v, vs) then $\gamma_{xv^{-1}w,yv^{-1}w}$ includes edge (w, ws), and vice-versa. Hence,

$$\begin{aligned}
\rho_e &= \max_{s\in S}\frac{1}{|G|}\sum_{g\in G}\frac{1}{\pi(g)\mathsf{P}(g,gs)}\sum_{\substack{x,y\in G\,|\,x\neq y,\\ (g,gs)\in\gamma_{xy}}}\pi(x)\pi(y)\\
&\le \max_{s\in S}\frac{1}{p(s)}\sum_{x,y\in G\,|\,x\neq y}N(x^{-1}y,s)\pi(x)\pi(y)\\
&< \max_{s\in S}\frac{1}{p(s)}\max_{g\in G}N(g,s)\,.
\end{aligned}$$

□